arXiv:1605.05309v1 [stat.ME] 17 May 2016

Identification and Estimation of Causal Effects with Outcomes Truncated by Death

Linbo Wang[1], Xiao-Hua Zhou[2] and Thomas S. Richardson[2]

[1]Harvard School of Public Health, Boston, Massachusetts, U.S.A.

[2]University of Washington, Seattle, Washington, U.S.A.

Abstract

It is common that in medical studies, the outcome of interest is truncated by death, meaning that a subject had died before the outcome could be measured. In this case, restricted analysis among survivors may be subject to selection bias. It is hence of interest to estimate the survivor average causal effect (SACE), defined as the average causal effect among subjects who would survive under either exposure. In this paper, we consider the identification and estimation problems of the SACE. We propose to use a "substitution variable" in place of the latent membership in the always-survivor group. The identifiability conditions required for a "substitution variable" are conceptually similar to conditions for an instrumental variable. We show that the SACE is identifiable with use of such a "substitution variable," and propose novel model parameterizations for estimation of the SACE under our identification assumptions. Our approaches are illustrated via simulation studies and two data analyses.

Keywords: Causal inference; Instrumental variable; Model parameterization; Survivor average causal effect.

1 Introduction

In medical studies, researchers are often interested in evaluating risk factors for a non-mortality outcome such as memory decline. However, the non-mortality outcome may be truncated by death if some subjects die before the follow-up assessment, leaving their non-mortality outcomes to be undefined. For example, suppose we are interested in estimating the effect of smoking on memory decline in an aged population.

If a subject dies before the follow-up memory test is administered, then his/her memory score at the follow-up visit is undefined. Direct comparisons between smokers and non-smokers among observed survivors are subject to selection bias as nonsmokers are more likely to survive to the follow-up assessment (Rosenbaum, 1984; Robins and Greenland, 1992). More fundamentally, direct comparisons among observed survivors are not causally interpretable as they compare outcomes from different subpopulations at baseline (Rubin, 2006). Alternatively, Robins (1986, §12.2) proposed to estimate the average causal effect in the always-survivor group, the group of subjects who would survive if they choose to receive either exposure at baseline; see also Rubin (1998); Frangakis and Rubin (1999, 2002). The resulting estimand is termed the survivor average causal effect (SACE).

The SACE is only partially identifiable without further assumptions (Zhang and Rubin, 2003). Large sample bounds for the SACE have been derived under minimal assumptions (Zhang and Rubin, 2003; Imai, 2008; Long and Hudgens, 2013). In order to identify the SACE, a common strategy is to perform a sensitivity analysis by assuming a class of identifiability conditions indexed by an interpretable sensitivity parameter (Gilbert et al., 2003; Shepherd et al., 2006; Hayden et al., 2005; Egleston et al., 2007; Jemiai et al., 2007; Chiba and VanderWeele, 2011). Alternatively, identification of the SACE can be based on covariate information. For example, Tchetgen Tchetgen (2014) introduced a *variant* of the SACE to be identified when risk factors of survival are available in post-exposure follow-up. The resulting causal contrast is for a principal stratum defined in term of particular post-exposure risk factors. Alternatively, to identify the SACE in a randomized study setting, Ding et al. (2011) proposed a semiparametric identification approach based on a baseline variable whose distribution is informative of the membership of the always-survivor group. With this baseline variable, they showed that the SACE was identifiable under their assumptions. However, as pointed out by Tchetgen Tchetgen (2014), their assumption essentially requires that there are no common causes of the survival and outcome processes, which is very unlikely even in randomized studies.

In this article, we relax Ding et al. (2011)'s identification assumptions by employing more detailed covariate information. In comparison to previous works that attempt to identify the SACE with covariate information, our causal estimand is defined independently of the incorporated covariates. Furthermore, the proposed approach is applicable to both randomized trials and observational studies, *and* allows for measured common causes of the survival and outcome processes. We also discuss possible violations and alternative approaches to our identification model. To estimate the SACE in practice, we impose additional distributional assumptions on the *observed data*. This is challenging for two reasons. First, unlike the

standard observational study setting, in our identification framework the baseline covariates may be common causes between (any pair of) treatment, survival and outcome. As we explain later in Section 4, this unique role for baseline covariates make the standard propensity score methods inappropriate in our setting. Second, novel model parameterizations are needed so that our distributional assumptions are compatible with our identifiability assumptions.

The rest of the article is organized as follows. In Section 2, we define our parameter of interest, namely the SACE. We first introduce our identification assumptions in Section 3, and then propose model parameterizations for estimation of the SACE in Section 4. Finite sample performance of the proposed estimators is evaluated by simulation studies described in Section 5. In Section 6, we apply our method to compare two prostate cancer treatments with data from a Southwest Oncology Group (SWOG) Trial and estimate the effect of smoking on memory decline with data from the Health and Lifestyle Study (HALS). We end with a discussion in Section 7.

2 Data structure, notation and causal estimand

Consider a medical study with a single follow-up visit. Let Z be the exposure indicator and W denote observed covariates at baseline. We assume that each subject has two potential survival outcomes $S(1)$ and $S(0)$, defined as the survival status at the follow-up visit that would have been observed had the subject been exposed and unexposed, respectively. Similarly we let $Y(1)$ and $Y(0)$ denote the potential (non-mortality) outcomes under exposure and non-exposure, respectively. We assume that for $z = 0, 1$, $Y(z)$ takes real values only if $S(z) = 1$. We extend the definition of $Y(z)$ so that it takes constant value $*$ if $S(z) = 0$.

We use G to denote the survival type as defined in Table 1. One can see from Table 1 that there exists a one-to-one mapping between the survival type and the bivariate potential survivals $(S(1), S(0))$. As a result, G can be interpreted as an abbreviation for $(S(1), S(0))$.

We adopt the axiom of consistency such that the observed outcome Y satisfies $Y = ZY(1) + (1 - Z)Y(0)$ and the observed survival S satisfies $S = ZS(1) + (1 - Z)S(0)$. The observed samples $O_i = (Z_i, W_i, S_i, Y_i), i = 1, \ldots, N$ are independently drawn from an infinite super-population.

Throughout this article, we assume that there is no interference between study subjects regarding both the survival S and the non-mortality outcome Y, and there is only one version of exposure (SUTVA, Rubin (1980)).

Table 1: Patient survival types

$S(1)$	$S(0)$	Survival type	G	Description
1	1	always-survivor	LL	The subject always survives, regardless of exposure status.
1	0	protected	LD	The subject survives if exposed, but dies if not exposed.
0	1	harmed	DL	The subject dies if exposed, but survives if not exposed.
0	0	doomed	DD	The subject always dies, regardless of exposure status.

Rubin (2000) noted that the observed survivors in the exposed group are from a mixture of always-survivor and protected strata, while the observed survivors in the non-exposed group are from a mixture of always-survivor and harmed strata. As a result, direct comparisons between different exposure groups among observed survivors are not causally meaningful as these people are from different subpopulations at baseline. Instead, as always-survivor is the only stratum such that both $Y(1)$ and $Y(0)$ take real values, we define our causal estimand to be the average causal effect in the always-survivor stratum:

$$SACE = E[Y(1) - Y(0) \mid G = LL].$$

3 Identification of the SACE

3.1 The identification problem

Despite its causal interpretability, the SACE is in general not identifiable as it depends on the potential outcomes $Y(1), Y(0), S(1)$ and $S(0)$. Furthermore, it is not identifiable even under standard assumptions in the causal inference literature, which we introduce below.

A.1 (Monotonicity) $S_i(1) \geq S_i(0)$ for all $i = 1, \ldots, N$.

The monotonicity assumption may be plausible in some observational studies. For example, in studies evaluating the effect of smoking on memory decline, it is widely believed that smoking is bad for the overall health, and hence the overall survival. This assumption tends to be shaky in randomized clinical trials with acute diseases because (typically) a clinical trial would be unethical if the researchers believe that one treatment benefits survival *a priori*. To address this issue, we relax this assumption later in Section 3.4.

A.2 (S-Ignorability) The treatment assignment is independent of the potential survivals given the observed covariates W, denoted as $Z \perp\!\!\!\perp S(z) \mid W$.

A.3 (Y-Ignorability) The treatment assignment in the always-survivor stratum is marginally independent

of the potential outcomes given observed covariates W, denoted as $Z \perp\!\!\!\perp Y(z) \mid W, G = LL; z = 0, 1$.

A.2 and A.3 are similar to the (weakly) ignorable treatment assignment assumption (see e.g. Imbens, 2000). Under A.3, we have

$$E[Y(z) \mid G = LL] = \frac{E_W\left\{\mu_{z,LL,W}\pi_{LL|W}\right\}}{E_W\{\pi_{LL|W}\}}, \tag{1}$$

where $\mu_{z,g,W} = E[Y \mid Z = z, G = g, W], \pi_{g|W} = P(G = g \mid W)$. Under A.1 and A.2, $\pi_{LL|W}$ and $\pi_{LD|W}$ can be identified by

$$\pi_{LL|W} = P(S = 1 \mid Z = 0, W), \;\; \pi_{LD|W} = P(S = 1 \mid Z = 1, W) - P(S = 1 \mid Z = 0, W); \tag{2}$$

$\mu_{0,LL,W}$ can be identified by

$$\mu_{0,LL,W} = E[Y \mid Z = 0, S = 1, W]. \tag{3}$$

However, $\mu_{1,LL,W}$ is not identifiable from the observed data. In fact, even if the outcome Y has bounded support, one can only partially identify $\mu_{1,LL,W}$ from the equation

$$E[Y \mid Z = 1, S = 1, W] = \frac{\pi_{LL|W}}{P(S = 1 \mid Z = 1, W)}\mu_{1,LL,W} + \frac{\pi_{LD|W}}{P(S = 1 \mid Z = 1, W)}\mu_{1,LD,W}.$$

Therefore, A.1–A.3 alone are not sufficient for identifying the SACE.

3.2 Identifying the SACE using a "substitution variable"

To identify the SACE, without loss of generality, we assume that the baseline covariates W can be written as (X, A). We propose an identification framework in which the role for X is similar to a "confounder" whereas the role for A is similar to an "instrument." Specifically, we make the following assumptions on A and X:

A.4 (Exclusion restriction) $\quad A \perp\!\!\!\perp Y(1) \mid Z = 1, G, X.$

A.5 (Substitution relevance) $\quad A \not\perp\!\!\!\perp G \mid Z = 1, S = 1, X.$

We note that the survival type G is a (latent) baseline variable that satisfies A.4 and A.5. For this reason,

any variable satisfying A.4 and A.5 is called a *substitution variable* for G. The conditions for a substitution variable are similar to those for an instrumental variable (Angrist et al., 1996): A.4 is similar to the exclusion restriction assumption in an instrumental variable analysis in that they both capture the notion that A has no "direct effect" on Y, and A.5 is similar to the instrumental relevance assumption; in particular, they both require that A be a relevant variable. As we illustrate later in Section 3.3, even in a randomized study setting, the inclusion of covariate information X makes A.4 and A.5 more plausible.

The following Theorem 1 states that the SACE is identifiable with a substitution variable for the survival type.

Theorem 1: Under A.1–A.5, the SACE is identifiable.

3.3 Illustration under a NPSEM-IE framework

The key assumptions in our identification model are the conditional independence assumptions A.2–A.4. These assumptions are implied by a certain Non-Parametric Structural Equations Model with Independent Errors (NPSEM-IE) (Pearl, 2009; Richardson and Robins, 2013), where the NPSEM is given by

$$X = f_X(\epsilon_X), A = f_A(X, \epsilon_A), Z = f_Z(X, A, \epsilon_Z), S = f_S(X, A, Z, \epsilon_S), Y = f_Y(X, Z, S, \epsilon_Y) \quad (4)$$

and the error terms $\epsilon_X, \epsilon_A, \epsilon_Z, \epsilon_Y$ are jointly independent. See Appendix B for a proof. Note here the error terms can be interpreted as the set of one-step-ahead counterfactuals; for example, ϵ_Z can be interpreted as $\{Z(a, x), a \in \mathcal{A}, x \in \mathcal{X}\}$, where $\mathcal{A}$ and $\mathcal{X}$ are the sample spaces for A and X, respectively. Figure 1 gives the simplest causal diagram associated with the NPSEM (4). There is no directed edge from A to Y, encoding the notion that A has no direct effect on Y. We also remark that our assumptions A.2–A.4 also follow from (4) with dependent errors that allow for unmeasured confounding between nodes in Fig. 1. See Appendix C for examples.

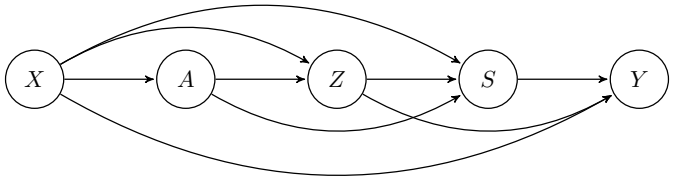

Figure 1: The simplest causal diagram associated with the NPSEM (4).

In comparison to the identification assumptions in Ding et al. (2011), we introduce baseline covariates X to the identification framework. Our approach is hence applicable to observational studies due to the presence of edges $X \rightarrow Z$, $X \rightarrow S$ and $X \rightarrow Y$. Moreover, we allow for edges $X \rightarrow S$ and $X \rightarrow Y$, thereby avoiding the assumption that there are no common causes of the survival and outcome processes, which is very restrictive even for randomized studies.

Remark 1: We note that the assumption of an NPSEM-IE cannot be tested, even in principle, via any randomized experiment on the measured variables (Robins and Richardson, 2010; Richardson and Robins, 2013).

3.4 Relaxing the monotonicity assumption

In this part we consider an alternative stochastic monotonicity assumption to relax the monotonicity assumption A.1. This alternative assumption has been used previously by Roy et al. (2008) in the context of non-compliance and Lee et al. (2010) in the context of truncation by death.

Specifically, we characterize the association between the potential survivals $S(1)$ and $S(0)$ by

$$\rho(W) = \frac{P(G = LL \mid W) - P(S(1) = 1 \mid W)P(S(0) = 1 \mid W)}{\min\{P(S(1) = 1 \mid W), P(S(0) = 1 \mid W)\} - P(S(1) = 1 \mid W)P(S(0) = 1 \mid W)}.$$

Note that $\rho(W) \leq 1$ and it differs from the correlation between $S(1)$ and $S(0)$ by a non-negative identifiable factor. In particular, the sign of $\rho(W)$ coincides with the sign of $Corr(S(1), S(0) \mid W)$. In practice one might be willing to assume that the potential survivals under active treatment and control are not negatively correlated such that

$$P(S(1) = 1 \mid S(0) = 1, W) \geq P(S(1) = 1 \mid S(0) = 0, W) \text{ almost surely.}$$

This gives rise to the following stochastic monotonicity assumption:

A.6 (Stochastic monotonicity)

$$0 \leq \rho(W) \leq 1 \text{ almost surely.}$$

We note that $\rho(W) = 0$ if and only if conditional on observed covariates W, $S(1)$ is not correlated with

$S(0)$. On the other hand, $\rho(W) = 1$ if and only if $P(G = LL \mid W) = \min\{P(S(1) = 1 \mid W), P(S(0) = 1 \mid W)\}$. Note A.6 does not reduce to A.1 when $\rho(W) = 1$; in particular, A.6 does not specify which treatment is more beneficial for survival. This is particularly appealing when the monotonicity assumption is not plausible *a priori* due to ethical reasons.

Theorem 2 states that under A.6, if we assume that A.4 and A.5 also hold in the control arm ($Z = 0$), then the SACE is identifiable from the observed data. The additional conditional independence condition (i) in Theorem 2 also follows from the NPSEM (4) with independent errors.

Theorem 2: If we assume A.2–A.6, and the additional conditions that

(i) $A \perp\!\!\!\perp Y(0) \mid Z = 0, X, G = LL$,

(ii) $A \not\perp\!\!\!\perp G \mid X, Z = 0, S = 1$,

then the SACE is identifiable given $\rho(W)$.

3.5 Relaxing the exclusion restriction assumption

In practice, the ignorability assumptions A.2 and A.3 may be questionable if there are unmeasured confounders for the survival S or the (non-mortality) outcome Y. Similarly, the exclusion restriction assumption A.4 may be questionable if the set of observed covariates X is not rich enough to contain all common causes of the substitution variable and outcome processes, as well as all common causes of the survival and outcome processes. There has been extensive literature on relaxing the ignorability assumptions for causal inference in observational studies (see e.g. Rosenbaum and Rubin, 1983; Imbens, 2003; Ding and Vanderweele, 2014), which may be applied here with slight modification. Hence in this article, we focus on alternative assumptions to relax the exclusion restriction assumption.

Specifically, we propose an alternative no-interaction assumption A.7, which assumes that conditioning on the baseline variables X, neither Z nor G modifies the effect of A on Y.

A.7 (no interaction) For all x, there exist $a_1 \neq a_0$ such that

$$\begin{aligned} &E[Y \mid Z = 1, G = LL, X = x, A = a_1] - E[Y \mid Z = 1, G = LL, X = x, A = a_0] \\ &= E[Y \mid Z = 1, G = LD, X = x, A = a_1] - E[Y \mid Z = 1, G = LD, X = x, A = a_0] \\ &= E[Y \mid Z = 0, G = LL, X = x, A = a_1] - E[Y \mid Z = 0, G = LL, X = x, A = a_0]. \end{aligned}$$

In contrast, the exclusion restriction assumption A.4 implies that the first two comparisons are both 0.

A.7 is similar in spirit to the no-interaction assumption in causal mediation analysis (see e.g. Robins and Greenland, 1992). A.7 holds under a wide range of semiparametric models. For example, a sufficient condition for A.7 is the following additive model:

$$E[Y \mid Z, G, X, A] = \eta_1(Z, G, X) + \eta_2(X, A), \tag{5}$$

where $G = 1$ for the LL group and $G = 0$ for the LD group; η_1 and η_2 are two functions that are completely unspecified. However, A.7 may still hold even if (5) were not true. Another feature of A.7 is that it can be used to identify the SACE with binary substitution variables. In comparison, to identify the SACE in a randomized study under exclusion restrcition violation, Ding et al. (2011) assumed a parametric model on $E[Y \mid Z = 1, G, A]$; furthermore their approach requires that the substitution variable A has at least three categories or is continuous.

Theorem 3 states that if we replace the exclusion restriction assumption A.4 with A.7, then the SACE is still identifiable.

Theorem 3: Under assumptions A.1–A.3, A.5 and A.7, the SACE is identifiable.

4 Model parameterization

In the previous section we have shown that the SACE is identifiable under various identification assumptions. To estimate the SACE in practice, unless the covariates W only take a few discrete values, we need to impose additional distributional assumptions. We first note that the identification assumptions imply certain constraints on the observed data law. Proposition 4.1 summarizes these constraints, with the proof in Appendix D.

Proposition 4.1:

(I) The assumptions of Theorem 1 imply the following constraints on the law of (Z, S, Y, W):

$$P(S = 0 \mid Z = 1, X, A) \leq P(S = 1 \mid Z = 1, X, A); \tag{6}$$

$$\text{for all } x, E[Y \mid Z=1, S=1, X=x, A=a] \text{ is bounded as a function of } a; \tag{7}$$

$$\text{for all } x, \frac{P(S=1 \mid Z=0, X=x, A=a)}{P(S=1 \mid Z=1, X=x, A=a)} \text{ is not constant as a function of } a. \tag{8}$$

(II) The assumptions of Theorem 2 imply the following constraints on the law of (Z, S, Y, W): either (8) holds, or $\forall z, x, P(S=1 \mid Z=z, X=x, A=a)$ is not a constant of a.

(III) The assumptions of Theorem 3 imply (6) and that

$$\text{for all } x, E[Y \mid Z=1, S=1, X=x, A=a] - E[Y \mid Z=0, S=1, X=x, A=a]$$
$$\text{is bounded as a function of } a. \tag{9}$$

To ensure that the modeling assumptions are compatible with the constraints in Proposition 4.1, we avoid imposing distributional assumptions on the observed data directly. Instead, we directly make the following distributional assumptions on the law of $(Z, W, S(1), S(0), Y(1), Y(0))$ and check that they are compatible with the identification assumptions. For brevity here we only give models under the assumptions of Theorem 1. Parameterizations for estimating the SACE under alternative identification assumptions which are similar in spirit are described in Appendix E.

M.1 $E[Y(0) \mid Z=0, G=LL, X, A]$ is known up to a finite-dimensional parameter α_1; that is, $E[Y(0) \mid G=LL] = m_1(X, A; \alpha_1)$, where $m_1(\cdot, \cdot; \alpha_1)$ is a known function and α_1 is an unknown parameter. Specifically, for our simulations and data examples we consider

$$m_1(X, A; \alpha_1) = \alpha_{10} + X^T \alpha_{11} + A\alpha_{12}. \tag{10}$$

M.2 $E[Y(1) \mid Z=1, G, X, A]$ is known up to a finite-dimensional parameter α_2; that is, $E[Y(1) \mid Z=1, G=g, X, A] = m_2(X, G; \alpha_2)$, where G takes values in $\{LL, LD\}$, $m_2(\cdot, \cdot; \alpha_1)$ is a known function and α_2 is an unknown parameter. Note that due to A.4, $E[Y(1) \mid Z=1, G, X, A]$ does not depend on the value of A. Specifically, for our simulations and data examples we code LL to be 1, LD to be 0, and consider

$$m_2(X, G; \alpha_2) = \alpha_{20} + X^T \alpha_{21} + G\alpha_{23}. \tag{11}$$

M.3 $P(S(1) = 1 \mid X, A)$ is known up to a finite-dimensional parameter β_1; that is $P(S(1) = 1 \mid X, A) = \theta_1(X, A; \beta_1)$, where $\theta_1(\cdot, \cdot; \beta_1)$ is a known function and β_1 is an unknown parameter. Specifically, for our simulations and data examples we consider

$$\theta_1(X, A; \beta) = \text{expit}(\beta_{10} + X^T \beta_{11} + A\beta_{12}). \tag{12}$$

M.4 $P(S(0) = 1 \mid X, A)/P(S(1) = 1 \mid X, A)$ is known up to a finite-dimensional parameter γ; that is $P(S(0) = 1 \mid X, A)/P(S(1) = 1 \mid X, A) = \theta_{0/1}(X, A; \gamma)$, where $\theta_{0/1}(\cdot, \cdot; \gamma)$ is a known function and γ is an unknown parameter. Specifically, for our simulations and data examples we consider

$$\theta_{0/1}(X, A; \gamma) = \text{expit}(\gamma_0 + X^T \gamma_1 + A\gamma_2). \tag{13}$$

Rather than M.3 and M.4, some previous researchers impose distributional assumptions on $P(S(z) \mid X, A)$, $z = 0, 1$ directly (e.g. Lee et al., 2010). However, due to the monotonicity assumption, $P(S(0) = 1 \mid X, A)$ resides in the range $[0, P(S(1) = 1 \mid X, A)]$. Hence the model parameters for $P(S(1) = 1 \mid X, A)$ live in a constrained space, making estimation and asymptotic analysis difficult. To avoid such constraints, we re-parameterize our parameters as in M.3 and M.4.

To derive the maximum likelihood estimator, we note that M.1–M.4 correspond to the following modeling constraints on the observed data distribution:

$$\begin{aligned}
P(S = 1 \mid Z = 1, X, A) &= \theta_1(X, A; \beta_1), \\
P(S = 1 \mid Z = 0, X, A) &= \theta_1(X, A; \beta_1)\theta_{0/1}(X, A; \gamma), \\
E[Y \mid Z = 1, S = 1, X, A] &= \theta_{0/1}(X, A; \gamma)m_2(X, 1; \alpha_2) + (1 - \theta_{0/1}(X, A; \gamma))m_2(X, 0; \alpha_2), \\
E[Y \mid Z = 0, S = 1, X, A] &= m_1(X, A; \alpha_1).
\end{aligned}$$

It is easy to see that these models are compatible with the testable implications of A.1–A.5 described in Proposition 4.1. Ordinary least squares and maximum likelihood estimation may be used for the parameter estimation. The SACE can then be estimated using (1), in which $\mu_{z,LL,W}$ can be estimated using (10) and (11), $\pi_{LL|W}$ can be estimated using the product of (12) and (13), and the integration is taken with respect to the empirical distribution of W. Using standard M-estimation theory, one can show that the resulting

estimate of SACE is consistent and asymptotically normally distributed.

Remark 2: Theorem 1 implies a plug-in estimator by solving equations (2), (3), (14) and (15). Our estimation method is different from this plug-in estimator. This is particularly convenient when A is continuous, in which case A needs to be discretized for the plug-in estimator to be applicable (Ding et al., 2011).

Remark 3: An alternative popular approach to estimate causal effects is based on the propensity score $e(X)$, defined as the probability of assignment to exposure conditioning on baseline covariates. However, although the propensity score is sufficient for summarizing the effect of X on Z in the sense that $Z \perp\!\!\!\perp X \mid e(X)$, it is not sufficient for summarizing the effect of X on A. Thus unlike a standard causal effect estimation setting, the propensity score methods are not directly applicable for estimating the SACE under our identification assumptions.

5 Simulation studies

5.1 Settings following A.1–A.5

We first consider simulation settings in which data were generated according to Fig. 1. Specifically the baseline covariates $X = (X_1, X_2, X_3)$ are a combination of discrete and continuous variables: X_1 is a discrete variable taking value 1 or -1 with probability 1/2, and (X_2, X_3) follows a multivariate normal distribution $N(\mu, \Sigma)$, where

$$\mu = \begin{pmatrix} 1 \\ -1 \end{pmatrix} \text{ and } \Sigma = \begin{pmatrix} 1 & 0.5 \\ 0.5 & 1 \end{pmatrix}.$$

Conditional on X, the substitution variable A was generated from a Bernoulli distribution such that $P(A = 1 \mid X) = \text{expit}(X^T \boldsymbol{u})$, where $\boldsymbol{u} = (1, 1, 1)^T$. The exposure variable Z was generated following a logistic model: $P(Z = 1 \mid X, A) = \text{expit}(\delta_1 X \mathbf{1} + \delta_1 A)$, where δ_1 is a parameter taking value 0 or 1. The survival type G was generated from a multinomial distribution such that $P(G = LL \mid X, A) = \text{expit}(\gamma_0 + X^T\gamma_1 + A\gamma_2)\text{expit}(\beta_{10} + X^T\beta_{11} + A\beta_{12})$, $P(G = LD \mid X, A) = \big\{1 - \text{expit}(\gamma_0 + X^T\gamma_1 + A\gamma_2)\big\}\,\text{expit}(\beta_{10} + X^T\beta_{11} + A\beta_{12})$ and $P(G = DD \mid X, A) = 1 - \text{expit}(\beta_{10} + X^T\beta_{11} + A\beta_{12})$, where $(\beta_{10}, \beta_{11}, \beta_{12}) = (2, \delta_2, \delta_2, \delta_2, 1)$, $(\gamma_0, \gamma_1, \gamma_2) = (0, -3\delta_2, \delta_2, \delta_2, 1)$ and δ_2 is a parameter taking values 0 or 1. The potential outcome $Y(z)$ was generated from the following normal distributions: $Y(1) \mid G = LL, X, A \sim N(5 +$

Table 2: Bias and standard deviation (in parenthesis) for various estimating methods of the SACE under settings following A.1–A.5. $\delta_1 = 1$ and $\delta_2 = 1$ corresponds to presence of confounding and common cause of the survival and outcome processes, respectively. Results are based on 1000 simulated data sets

Sample size	δ_1	δ_2	Estimation method		
			Naive	DGYZ	Prop-ER
200	0	0	-0.73(0.14)	2.76(87.92)	0.01(1.54)
		1	-0.47(0.15)	-7.90(134.36)	-0.12(0.25)
	1	0	-0.80(0.19)	-1.29(0.52)	-0.11(2.22)
		1	-0.41(0.18)	-1.83(2.40)	-0.27(0.35)
500	0	0	-0.73(0.08)	0.34(3.73)	0.11(0.77)
		1	-0.47(0.09)	-3.43(215.02)	-0.05(0.17)
	1	0	-0.81(0.12)	-1.27(0.27)	0.11(1.19)
		1	-0.42(0.12)	-1.46(0.88)	-0.22(0.23)
1000	0	0	-0.73(0.06)	0.06(0.38)	0.05(0.37)
		1	-0.47(0.07)	-14.47(188.83)	-0.03(0.12)
	1	0	-0.81(0.08)	-1.26(0.18)	0.10(0.61)
		1	-0.41(0.08)	-1.42(0.57)	-0.11(0.23)

$\delta_2 X^T \boldsymbol{u}, 0.5^2)$, $Y(1) \mid G = LD, X, A \sim N(3 + \delta_2 X^T \boldsymbol{u}, 0.5^2)$ and $Y(0) \mid G = LL, X, A \sim N(4 + \delta_2 X^T \boldsymbol{u}, 0.5^2)$. The observed survival S and observed outcome Y follows from the consistency assumption. Note under our settings, the treatment assignment is confounded when $\delta_1 = 1$, and X is a common cause of the survival and outcome processes when $\delta_2 = 1$. The true value for the SACE is 1.

We compare three methods for estimating the SACE: linear regression of Y on Z, X and A among observed survivors (Naive); the estimation method of Ding et al. (2011) (DGYZ); the proposed method under the exclusion restriction assumption A.4, i.e. the estimation method using M.1–M.4 (Prop-ER). Note that Prop-ER differs from DGYZ in both the identification conditions and estimation method. Table 2 summarizes the results. The naive regression method is biased in all settings due to selection bias. In the presence of confounding or common causes of S and Y ($\delta_1 \neq 0$ or $\delta_2 \neq 0$), DGYZ is inconsistent. However, even in the absence of confounding or common causes of S and Y ($\delta_1 = \delta_2 = 0$), DGYZ can be unstable when the sample size is small to moderate. Although in this case the substitution relevance assumption of Ding et al. (2011) holds such that $P(A = 1 \mid G = LL, Z = 1) \neq P(A = 1 \mid G = LD, Z = 1)$, for some simulated samples the estimate of $P(A = 1 \mid G = LL, Z = 1)$ can be very close to the estimate of $P(A = 1 \mid G = LD, Z = 1)$, leading to instability of the causal effect estimates. In contrast, Prop-ER performs well in all settings considered here.

Table 3: Bias and standard deviation (in parenthesis) for various estimating methods of the SACE under settings with exclusion restriction violation. $\delta_1 = 1$ and $\delta_2 = 1$ corresponds to presence of confounding and common cause of the survival and outcome processes, respectively. Results are based on 1000 simulated data sets

Sample size	δ_1	δ_2	Estimation method		
			DGYZ	Prop-ER	Prop-NI
200	0	0	7.44(180.19)	1.48(2.53)	0.36(0.42)
		1	-11.26(247.45)	0.19(0.38)	0.04(0.20)
	1	0	0.34(0.59)	1.34(2.70)	0.47(0.54)
		1	-1.25(3.25)	0.54(0.42)	0.27(0.29)
500	0	0	2.00(6.01)	1.68(0.95)	0.14(0.33)
		1	-1.17(355.16)	0.10(0.22)	0.02(0.13)
	1	0	0.43(0.29)	1.64(1.19)	0.25(0.38)
		1	-0.71(1.35)	0.41(0.37)	0.16(0.19)
1000	0	0	1.62(0.42)	1.62(0.44)	0.08(0.25)
		1	-20.78(285.28)	0.08(0.14)	0.02(0.09)
	1	0	0.43(0.20)	1.67(0.95)	0.11(0.29)
		1	-0.62(0.84)	0.31(0.50)	0.10(0.14)

5.2 Settings with exclusion restriction violation

We now evaluate the sensitivity of the proposed method to departure from the exclusion restriction assumption A.4. We generated data in the same way as in Section 5.1, except that the potential outcome $Y(z)$ was generated from the following normal distributions: $Y(1) \mid G = LL, X, A \sim N(5 + \delta_2 X^T \boldsymbol{u} + A, 0.5^2)$, $Y(1) \mid G = LD, X, A \sim N(3+\delta_2 X^T \boldsymbol{u}+A, 0.5^2)$ and $Y(0) \mid G = LL, X, A \sim N(4+\delta_2 X^T \boldsymbol{u}+A, 0.5^2)$. We also implemented the method Prop-NI, the proposed method under the no-interaction assumption A.7; that is, the estimation method using M.3, M.4 and M.7. The method of Ding et al. (2011) for dealing with exclusion restriction violations is not applicable here as they require the substitution variable A to be either continuous or have at least three categories.

The simulation results are presented in Table 3. As expected, Prop-NI is consistent for all simulation settings considered here, whereas both DGYZ and Prop-ER are biased for estimating the SACE. However, compared with DGYZ, the inclusion of covariate information X makes Prop-ER more robust to departure from the exclusion restriction assumption; specifically under small to moderate sample sizes, this advantage remains even with $\delta_1 = \delta_2 = 0$.

6 Data analysis

6.1 Application to a SWOG trial

We illustrate the advantage of the proposed estimation method using data from a randomized phase III trial to compare docetaxel plus estramustine with mitoxantrone plus prednisone in men with metastatic, hormone independent prostate cancer (Petrylak et al., 2004). The data set we use was created by Ding et al. (2011), which contains observations on 487 men aged from 47 to 88. Of these subjects, 258 were assigned to receive docetaxel plus estramustine ($Z = 1$) and 229 were assigned to receive mitoxantrone plus prednisone ($Z = 0$). In our analysis, we are interested in comparing these two treatments in term of health related quality of life (QoL) one year after receiving the treatment.

A naive analysis shows that among patients who survived at one year after receiving the assigned treatment, the QoL for those assigned to the docetaxel plus estramustine group is higher by 2.46 units (95% CI = [-3.31,8.24]) compared to those assigned to the mitoxantrone plus prednisone group. However, this estimate is not causally interpretable as subjects who would survive if assigned to docetaxel plus estramustine are different from subjects who would survive if assigned to mitoxantrone plus prednisone. Moreover, as reported by Petrylak et al. (2004), docetaxel plus estramustine is beneficial for the overall survival compared to mitoxantrone plus prednisone. The direct comparison among observed survivors is hence also subject to selection bias. Instead, we apply the following three methods that take account of truncation due to death: DGYZ, Prop-ER and Prop-NI. To account for possible common causes of the survival and outcome processes, we adjust for the following variables in the models used by Prop-ER and Prop-NI: age, race (black vs non-black), type of prognosis (PSA only vs others), bone pain (grade < 2 vs grade ≥ 2) and performance status ($0 - 1$ vs $2 - 3$). Following Ding et al. (2011), we use the baseline QoL as our substitution variable, and the change in QoL in the one-year period as our outcome.

We first analyze this data set under the monotonicity assumption. The point estimates as well as bootstrapped standard errors are displayed in Table 4. DGYZ suggested that docetaxel plus estramustine had a significant causal effect on the QoL among those who would survive one year after receipt of treatment regardless of which treatment group they were assigned to. In contrast, after accounting for the baseline covariate information that might simultaneously impact the (potential) survival and the (potential) QoL, we were not able to reach such a conclusion. Both Prop-ER and Prop-NI yield a point estimate that is much closer to 0, and their 95% confidence intervals cover 0. These results show that in order to

Table 4: Survivor average causal effect of docetaxel plus estramustine on health related quality of life

Estimation method	Point estimate	Bootstrapped SE	2.5%	50%	97.5%
DGYZ†	7.01	3.09	1.81	6.56	13.64
Prop-ER	3.06	11.79	-15.15	3.29	22.60
Prop-NI	2.73	3.82	-3.97	2.95	10.83

†: Results adapted from Ding et al. (2011)

control confounding between S and Y, it is important to incorporate baseline covariate information even in randomized trials.

We also note that Prop-ER is very unstable under the monotonicity assumption. In particular, the point estimate for the model parameter γ in M.4 is large in magnitude, suggesting that the true value of γ may lie on the boundary. This indicates a monotonicity violation. In fact, previous analyses also suggest that the monotonicity assumption might be problematic for this data set (Ding et al., 2011) and as we discussed earlier, for randomized trials the monotonicity assumption is not plausible *a priori*. For these reasons, although the overall one year survival rate in the docetaxel plus estramustine group (49.6%) is higher than that in the mitoxantrone plus prednisone group (38.9%), a sensitivity analysis of the monotonicity assumption is warranted. In our analysis the sensitivity parameter is π_{DL}, which has a one-to-one correspondence with the conventional sensitivity parameter ρ but has the interpretation that it is the overall fraction of patients who would die one year after receiving docetaxel plus estramustine but would survive one year after receiving mitoxantrone plus prednisone. Under our modeling assumptions M.2, M.5 and M.6, the range for π_{DL} is estimated to be [0.01,0.19], whose lower limit corresponds to $\rho = 1$. As the lower limit is greater than 0, the monotonicity assumption A.1 is not compatible with these modeling assumptions. On the other hand, the estimate of π_{DL} from Ding et al. (2011) is much greater than the upper limit, corresponding to a negative correlation between the potential survival under docetaxel plus estramustine and the potential survival under mitoxantrone plus prednisone. Figure 2 summarizes the results from our sensitivity analysis. As both of the sensitivity lines cross 0, without prior knowledge on the possible values for π_{DL}, we cannot reach any definitive answer about the causal effect of docetaxel plus estramustine versus mitoxantrone plus prednisone on the QoL one year after receiving the treatment. We can also see that the curves estimated by Prop-ER and Prop-NI agree approximately unless π_{DL} is smaller than 0.05.

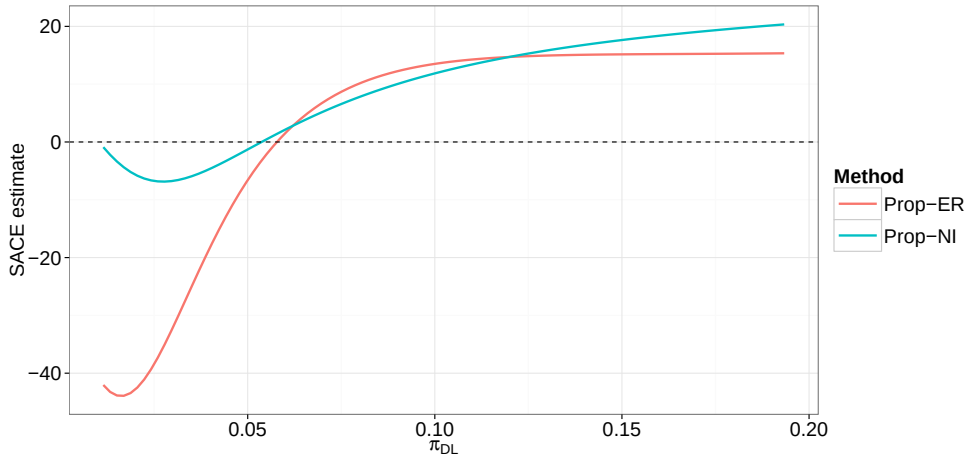

Figure 2: Sensitivity analysis for estimating the survivor average causal effect in the SWOG dataset.

6.2 Application to the Health and Lifestyle Study

We now use the proposed methods to estimate the effect of smoking on memory decline using observational data from the Health and Lifestyle Study (HALS). Smoking is a leading cause of preventable disease and death. There has been consensus that smoking is harmful to the overall health of smokers, and leads to increased risks of lung cancer, heart disease and stroke. However, there have been mixed findings about effects of smoking on cognitive functioning (Anstey et al., 2007; Peters et al., 2008). Some researchers have conjectured that this inconsistency may be associated with selection bias due to censoring by death (Hernán et al., 2008).

To investigate the impact of this selection bias, we use data from the Health and Lifestyle Survey (HALS), a population-based prospective cohort study conducted in England, Scotland and Wales (Cox et al., 1987). There were two visits in the HALS. The baseline survey HALS1 was conducted in 1984-1985 and interviewed 9003 participants on their health statuses, attitudes to health and other measurements related to health and lifestyles. Among these participants, 7414 agreed to have a nurse-visit where an incidental memory test was conducted. In this test, each respondent got a score ranging from 0 to 10 based on how many items he/she correctly recalled from a memory test. The follow-up survey HALS2 was conducted in 1991-1992, at which time 808 of the 7414 original respondents to the baseline survey had died. The surviving respondents of HALS1 were re-surveyed to examine their changes on memory scores over the 7-year period.

In our analysis, we use Z to denote the smoking status at the baseline visit, where $Z = 1$ for ever

smokers. We have $S = 1$ for surviving respondents from HALS1, and $S = 0$ for respondents of HALS1 who died before HALS2 was administered. The covariates X included age, sex, education level, alcohol consumption, BMI, household income and family smoking history. This is a superset of the potential confounders considered in Whittington and Huppert (1997), who studies the same problem with HALS using a restricted analysis to observed survivors. Using the naive method, but adjusting for observed confounders by classical linear regression smokers are estimated to have 0.04 more points of decline in memory score compared to non-smokers (95% CI = [-0.12,0.21]). The result is non-significant, suggesting that no evidence is found for effects of smoking on memory decline. However, it is well-known that smoking is bad for survival. An analysis using a logistic regression model shows that after adjusting for X, smoking is associated with a 94.8% (95% CI = [27.4%, 198.1%]) decrease in odds of surviving between the follow-ups in the HALS. This provides strong evidence for selection bias resulting from a naive analysis.

We then apply Prop-ER and Prop-NI to this data set. We assume that smoking does not save lives such that the monotonicity assumption holds. Family history of lung cancer is used as our substitution variable for the survival type as it is plausible that family history of lung cancer is not related to memory decline conditioning on the smoking behavior of the respondent and his/her parents (exclusion restriction assumption). On the other hand, as family history of lung cancer is indicative of "lung cancer genes" that are likely to be of low frequency but high penetrance (Satcher et al., 2002), it is likely to affect smoking-related mortality (substitution relevance).

Table 5 summarizes results from the proposed methods. The point estimates with Prop-ER and Prop-NI are very close to the estimate from the naive analysis, suggesting that the selection bias due to death is minimal for this data set. Although the estimates with different methods are close, it should be noted that they have different interpretations. The proposed methods estimate the SACE; that is, the effect of smoking on memory decline among respondents who would survive at the time of HALS2 regardless of whether they ever smoked at the time of HALS1. As this group of people are defined at baseline, the SACE is causally interpretable. In contrast, the naive method compares the memory decline of smokers to non-smokers among respondents who were observed to survive at the time of HALS2. This comparison cannot be interpreted causally as observed survivors at the time of HALS2 represent different groups of respondents at the baseline visit.

Table 5: Summary of results from the proposed methods

Method	Point estimate	Bootstrapped SE	95% CI
Naive	0.043	0.084	[-0.121,0.207]
Prop-ER	0.007	0.099	[-0.153,0.244]
Prop-NI	0.022	0.088	[-0.133,0.205]

7 Discussion

In this article, we have considered the identification and estimation problem of the SACE. Compared with previous works that attempt to identify the SACE using an exogenous or instrumental variable, we introduce baseline covariates X into the identification framework, which may be common causes between (any pair of) treatment, survival and outcome. This role for baseline covariates is unique to our identification framework for the SACE. As a result, in contrast to the standard estimation problem of the average causal effect (ACE), inclusion of baseline covariates is crucial for obtaining unbiased estimates of the SACE even in randomized studies when, as is likely, there may be confounding between survival and the post-survival outcome.

In our estimation approach, instead of imposing modeling assumptions on the observed data, we use parametric models for the potential outcomes. This not only simplifies the estimation procedure, but also makes it easier to ensure compatibility between our identifiability and modeling assumptions. We also propose alternative methods to relax the monotonicity assumption and the exclusion restriction assumption. Formal inference under our alternative assumption A.6 requires a sensitivity analysis framework that accounts for uncertainty due to both (partial) non-identifiability and sampling variability, which can be based on ignorance and uncertainty regions as discussed by Vansteelandt et al. (2006) and others. To focus on the main ideas here we only present point estimates under A.6 for fixed values of ρ.

In the current work, we have only considered the binary exposure variable case. In practical medical studies, there may be multiple levels of the exposure variable. For example, the smoking variable may be coded as never smokers, past smokers and current smokers. In this case, it would be interesting to generalize the proposed methods to deal with ordinal exposure status.

Acknowledgements

The authors thank Peng Ding, Peter Gilbert, James Robins and Eric Tchetgen Tchetgen for valuable discussions. The authors also thank members of National Alzheimer's Coordinating Center (U01 AG016976) and MELODEM selection group for their helpful comments. This work was initiated when the first author was a graduate student at the University of Washington. Zhou was supported in part by a US Department of Veterans Affairs, Veterans Affairs Health Administration, Research Career Scientist award (RCS 05-196). Richardson was supported by National Institutes of Health Grant R01 AI032475. Zhou is also affiliated with Health Services Research and Development Center of Excellence, Veterans Affairs Puget Sound Health Care System, Seattle.

Appendix

A Proof of Theorem 1–3

By definition, $SACE = E[Y(1) \mid G = LL] - E[Y(0) \mid G = LL]$, where $E[Y(z) \mid G = LL]$ is given in (1). It then suffices to identify $\pi_{LL|W}$ and $\mu_{z,LL,W}$ for identification of the SACE.

Identification of $\pi_{LL|W}$: under the assumptions of Theorem 1 or 3, $\pi_{LL|W}$ can be identified from (2), whereas under the assumptions of Theorem 2, $\pi_{LL|W}$ can be identified from the following equations:

$$
\begin{aligned}
P(S(1) = 1 \mid W) &= P(S = 1 \mid Z = 1, W) = \pi_{LL|W} + \pi_{LD|W};\\
P(S(0) = 1 \mid W) &= P(S = 1 \mid Z = 0, W) = \pi_{LL|W} + \pi_{DL|W};\\
\pi_{LL|W} &= P(S(0) = 1 \mid W) P(S(1) = 1 \mid S(0) = 1, W);\\
P(S(1) = 1 \mid S(0) = 1, W) &= P(S(1) = 1 \mid W) + \rho(W) \left(\min \left\{ 1, \frac{P(S(1) = 1 \mid W)}{P(S(0) = 1 \mid W)} \right\} - P(S(1) = 1 \mid W) \right).
\end{aligned}
$$

Note under these assumptions, $\pi_{g|W}$ is also identified for $g \in \{LD, DL, DD\}$.

Identification of $\mu_{z,LL,W}$: under the assumptions of Theorem 1, $\mu_{0,LL,W}$ can be identified from (3). For $z = 1$ and two distinct values a_1 and a_0 that lie in the range of A,

$$
E[Y \mid Z = z, G = g, X = x, A = a_1] = E[Y \mid Z = z, G = g, X = x, A = a_0]. \tag{14}
$$

Hence $\mu_{z,LL,w} = E[Y \mid Z = z, G = g, X = x] \equiv \mu_{z,LL,x}$. Furthermore, $\mu_{z,LL,x}$ can be identified from

$$\begin{aligned} E[Y \mid Z = z, S = 1, X = x, A = a_1] &= p_{LL|z,x,a_1}\mu_{z,LL,x} + p_{LD|z,x,a_1}\mu_{z,LD,x}, \\ E[Y \mid Z = z, S = 1, X = x, A = a_0] &= p_{LL|z,x,a_0}\mu_{z,LL,x} + p_{LD|z,x,a_0}\mu_{z,LD,x}, \end{aligned} \tag{15}$$

where for $g = LL$ or LD, $p_{g|z,x,a} \equiv P(G = g \mid X = x, A = a)/P(S = 1 \mid Z = z, X = x, A = a)$ is identifiable from the observed data as described above.

Under the assumptions of Theorem 2, we can obtain similarly that for $z \in \{0, 1\}$, $\mu_{z,LL,w} = \mu_{z,LL,x}$, where $\mu_{z,LL,x}$ can be identified from (15).

Under the assumptions of Theorem 3, similarly we get that $\mu_{0,LL,W}$ can be identified from (3) and $\mu_{1,LL,w} = \mu_{1,LL,x}$, where $\mu_{1,LL,x}$ can be identified from the following equations:

$$\begin{aligned} E[Y \mid Z = 1, S = 1, X = x, A = a_1] &= p_{LL|1,x,a_1}\mu_{1,LL,x,a_1} + p_{LD|1,x,a_1}\mu_{1,LD,x,a_1}; \\ E[Y \mid Z = 1, S = 1, X = x, A = a_0] &= p_{LL|1,x,a_0}\mu_{1,LL,x,a_0} + p_{LD|1,x,a_0}\mu_{1,LD,x,a_0}; \\ \mu_{1,LL,x,a_1} - \mu_{1,LL,x,a_0} &= \mu_{1,LD,x,a_1} - \mu_{1,LD,x,a_0}; \\ \mu_{1,LL,x,a_1} - \mu_{1,LL,x,a_0} &= \mu_{0,LL,x,a_1} - \mu_{0,LL,x,a_0}; \\ E[Y \mid Z = 0, S = 1, X = x, A = a_1] &= \mu_{0,LL,x,a_1}; \\ E[Y \mid Z = 0, S = 1, X = x, A = a_0] &= \mu_{0,LL,x,a_0}; \\ \mu_{1,LL,x} &= p_{a_1|1,LL,x}\mu_{1,LL,x,a_1} + p_{a_0|1,LL,x}\mu_{1,LL,x,a_0}, \end{aligned}$$

in which $\mu_{z,g,x,a} = E[Y \mid Z = z, G = g, X = x, A = a]$ and $p_{a|1,LL,x} = P(A = a \mid Z = z, G = g, X = x)$ can be identified as follows:

$$\begin{aligned} P(A = a \mid Z = z, G = g, X = x) &= \frac{P(G = g \mid Z = z, X = x, A = a)P(A = a \mid Z = z, X = x)}{P(G = g \mid Z = z, X = x)} \\ &= \frac{\pi_{g|w}P(A = a \mid Z = z, X = x)}{\sum_a \pi_{g|w}P(A = a \mid Z = z, X = x)}. \end{aligned}$$

B Proof that the NPSEM-IE in Section 3.3 implies A.2–A.4

Under (4), $\epsilon_Z \perp\!\!\!\perp (\epsilon_X, \epsilon_A, \epsilon_S, \epsilon_Y)$ implies that $Z \perp\!\!\!\perp (\epsilon_S, \epsilon_Y) \mid X, A$, which then implies A.2 and A.3; recall here that the error terms can be interpreted as the set of one-step-ahead counterfactuals. To see that A.4

holds under the NPSEM (4), note the following:

$$
\begin{aligned}
&P(Y|A=a_1, Z=1, X=x, S(1)=1, S(0)) \\
&= P(Y|A=a_1, Z=1, X=x, S=1, S(0)) \quad \text{(consistency)} \\
&= P(Y|A=a_1, Z=1, X=x, S=1) \quad (Y \perp\!\!\!\perp S(0) \mid A, Z=1, X, S=1) \\
&= P(Y|A=a_0, Z=1, X=x, S=1) \quad (Y \text{ is d-separated from } A \text{ conditioning on } Z, X, S) \\
&= P(Y|A=a_0, Z=1, X=x, S=1, S(0)) \quad (Y \perp\!\!\!\perp S(0) \mid A, Z=1, X, S=1) \\
&= P(Y|A=a_0, Z=1, X=x, S(1)=1, S(0)) \quad \text{(consistency)},
\end{aligned}
$$

where $Y \perp\!\!\!\perp S(0) \mid A, Z=1, X, S=1$ as Y is d-separated from ϵ_S given X, A, Z, S.

As A.4 holds trivially for the subgroup with $S(1)=1$, we complete the proof. □

C Alternative causal diagrams compatible with the NPSEM (4)

Figure 3 provide some more complicated causal diagrams that imply assumptions A.2–A.4 but allow for dependence of errors in the NPSEM (4) representing unmeasured confounding between nodes in Fig. 1. Here unmeasured confounding is denoted as a bi-directed edge between observed nodes. In general, using graphical terminology (Richardson, 2003) the mixed graphs (such as the ones in Fig. 3) imply A.2–A.4 as long as the pairs (Z, S), (Z, Y), (A, Y) and (S, Y) are not in the same "district," where a district is a connected component of the graph obtained by removing all edges that are not bi-directed. Specifically, A.2 requires that Z and S cannot be in the same district, A.3 requires that neither Z and Y nor S and Y can be in the same district, and A.4 requires that neither A and Y nor S and Y can be in the same district.

D Proof of Proposition 4.1

The only non-standard constraints are (7) and (9). To prove the former, note that under A.4, (15) holds for all a and $z=1$, i.e.

$$E[Y \mid Z=1, S=1, X=x, A=a] = p_{LL|x,a}\mu_{1,LL,x} + p_{LD|x,a}\mu_{1,LD,x}. \tag{16}$$

It follows that

$$|E[Y \mid Z=1, S=1, X=x, A=a]| \leq \max\{|\mu_{1,LL,x}|, |\mu_{1,LD,x}|\}.$$

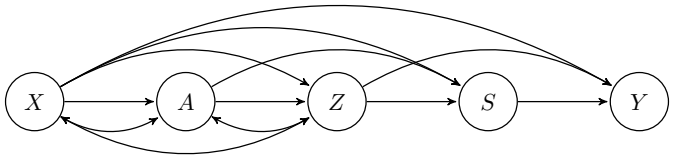

(i)

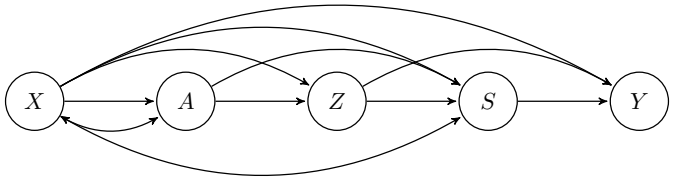

(ii)

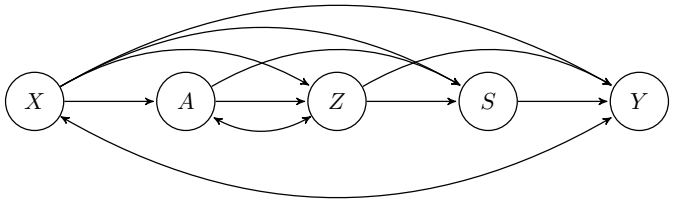

(iii)

Figure 3: More complicated causal diagrams that imply A.2–A.4. The districts in these graphs are: (i) $\{X, A, Z\}, \{S\}$ and $\{Y\}$; (ii) $\{X, A, S\}, \{Z\}$ and $\{Y\}$; (iii) $\{X, Y\}, \{A, Z\}$ and $\{S\}$.

Hence for all x, $E[Y \mid Z = 1, S = 1, X = x, A = a]$ is bounded as a function of a. On the other hand, suppose that for all x, $E[Y \mid Z = 1, S = 1, X = x, A = a]$ is bounded. Let

$$\bar{f}(x) = \max_a E[Y \mid Z = 1, S = 1, X = x, A = a], \quad \text{and} \quad \underline{f}(x) = \min_a E[Y \mid Z = 1, S = 1, X = x, A = a].$$

Then (16) holds with

$$p_{LL|x,a} = \frac{\bar{f}(x) - E[Y \mid Z = 1, S = 1, X = x, A = a]}{\bar{f}(x) - \underline{f}(x)}, \quad \mu_{1,LL,x} = \underline{f}(x) \quad \text{and} \quad \mu_{1,LD,x} = \bar{f}(x).$$

Hence (7) summarizes all the constraints on the observed data law derived from (16).

The proof of (9) follows similarly and is hence omitted.

E Parameterizations under alternative identification assumptions

To estimate the SACE under the assumptions of Theorem 2, we assume M.2, M.3 and the following model:

M.5 $E[Y(0) \mid Z = 0, G, X, A]$ is known up to a finite-dimensional parameter α_3; that is, $E[Y(0) \mid Z = 0, G, X, A] = m_3(X, G; \alpha_3)$, where G takes values in $\{LL, DL\}$, $m_3(\cdot, \cdot; \alpha_3)$ is a known function and α_3 is an unknown parameter. Note that due to condition (i) of Theorem 2, $m_3(\cdot, \cdot; \alpha_3)$ does not depend on the value of A. Specifically, for our simulations and data examples we code LL to be 1, DL to be 0, and consider

$$m_3(X, G; \alpha_3) = \alpha_{30} + X^T \alpha_{31} + G\alpha_{33}.$$

M.6 $P(S(0) = 1 \mid X, A)$ is known up to a finite-dimensional parameter β_0; that is $P(S(0) = 1 \mid X, A) = \theta_0(X, A; \beta_0)$, where $\theta_0(\cdot, \cdot; \beta_0)$ is a known function and β_0 is an unknown parameter. Specifically, for our simulations and data examples we consider

$$\theta_0(X, A; \beta_0) = \text{expit}(\beta_{00} + X^T \beta_{01} + A\beta_{02}).$$

Similarly, to estimate the SACE under the assumptions of Theorem 3, we assume M.3, M.4 and the following model:

M.7 $E[Y \mid Z, G, X, A], g = LL, LD$ is known up to a finite-dimensional parameter α_4; that is, $E[Y \mid Z, G, X, A] = m_4(X, A, G, Z; \alpha_4)$, where G takes values in $\{LL, LD\}$, $m_4(\cdot, \cdot, \cdot, \cdot; \alpha_4)$ is a known function and α_4 is an unknown parameter. Note that due to A.7, $m_4(X, A, G, Z; \alpha_4)$ should not contain interaction terms within the pairs (A, Z) or (A, G). Specifically, for our simulations and data examples we code LL to be 1, DL to be 0 and consider

$$m_4(X, A, G, Z; \alpha_4) = \alpha_{40} + X^T\alpha_{41} + A\alpha_{42} + G\alpha_{43} + Z\alpha_{44}.$$

Finally, if one wishes to relax both the monotonicity assumption A.1 and the exclusion restriction A.4, one may assume M.3, M.6, and the following model:

M.8 $E[Y \mid Z, G, X, A]$ is known up to a finite-dimensional parameter α_5; that is, $E[Y \mid Z, G, X, A] = m_5(X, A, G, Z; \alpha_5)$, where G takes values in $\{LL, LD, DL\}$, $m_5(\cdot, \cdot, \cdot, \cdot; \alpha_5)$ is a known function and α_5 is an unknown parameter. Note that due to A.7, $m_5(X, A, G, Z; \alpha_5)$ should not contain interaction terms within the pairs (A, Z) or (A, G). Specifically, for our simulations and data examples we code LL to be 1, DL and LD to be 0 and consider

$$E[Y \mid Z, G, X, A] = \alpha_{50} + X^T\alpha_{51} + A\alpha_{52} + ZG\alpha_{53} + Z\alpha_{54} + (1 - Z)G\alpha_{55}.$$

References

Angrist, J. D., Imbens, G. W., and Rubin, D. B. (1996). Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91(434):444–455.

Anstey, K. J., von Sanden, C., Salim, A., and O'Kearney, R. (2007). Smoking as a risk factor for dementia and cognitive decline: a meta-analysis of prospective studies. *American Journal of Epidemiology*, 166(4):367–378.

Chiba, Y. and VanderWeele, T. J. (2011). A simple method for principal strata effects when the outcome has been truncated due to death. *American Journal of Epidemiology*, 173(7):745–751.

Cox, B., Blaxter, M., Buckle, A., Fenner, N., Golding, J., Gore, M., Huppert, F., Nickson, J., Roth, S. M., Stark, J., et al. (1987). *The Health and Lifestyle Survey. Preliminary report of a nationwide survey of the*

physical and mental health, attitudes and lifestyle of a random sample of 9,003 British adults. Health Promotion Research Trust.

Ding, P., Geng, Z., Yan, W., and Zhou, X.-H. (2011). Identifiability and estimation of causal effects by principal stratification with outcomes truncated by death. *Journal of the American Statistical Association*, 106(496):1578–1591.

Ding, P. and Vanderweele, T. J. (2014). Generalized cornfield conditions for the risk difference. *Biometrika*, 101(4):971–977.

Egleston, B. L., Scharfstein, D. O., Freeman, E. E., and West, S. K. (2007). Causal inference for non-mortality outcomes in the presence of death. *Biostatistics*, 8(3):526–545.

Frangakis, C. E. and Rubin, D. B. (1999). Addressing complications of intention-to-treat analysis in the combined presence of all-or-none treatment-noncompliance and subsequent missing outcomes. *Biometrika*, 86(2):365–379.

Frangakis, C. E. and Rubin, D. B. (2002). Principal stratification in causal inference. *Biometrics*, 58(1):21–29.

Gilbert, P. B., Bosch, R. J., and Hudgens, M. G. (2003). Sensitivity analysis for the assessment of causal vaccine effects on viral load in HIV vaccine trials. *Biometrics*, 59(3):531–541.

Hayden, D., Pauler, D. K., and Schoenfeld, D. (2005). An estimator for treatment comparisons among survivors in randomized trials. *Biometrics*, 61(1):305–310.

Hernán, M. A., Alonso, A., and Logroscino, G. (2008). Cigarette smoking and dementia: potential selection bias in the elderly. *Epidemiology*, 19(3):448–450.

Imai, K. (2008). Sharp bounds on the causal effects in randomized experiments with "truncation-by-death". *Statistics & Probability Letters*, 78(2):144–149.

Imbens, G. W. (2000). The role of the propensity score in estimating dose-response functions. *Biometrika*, 87(3):706–710.

Imbens, G. W. (2003). Sensitivity to exogeneity assumptions in program evaluation. *American Economic Review*, 93(2):126–132.

Jemiai, Y., Rotnitzky, A., Shepherd, B. E., and Gilbert, P. B. (2007). Semiparametric estimation of treatment effects given base-line covariates on an outcome measured after a post-randomization event occurs. *Journal of the Royal Statistical Society, Series B: Methodological*, 69(5):879–901.

Lee, K., Daniels, M. J., and Sargent, D. J. (2010). Causal effects of treatments for informative missing data due to progression/death. *Journal of the American Statistical Association*, 105(491):912–929.

Long, D. M. and Hudgens, M. G. (2013). Sharpening bounds on principal effects with covariates. *Biometrics*, 69(4):812–819.

Pearl, J. (2009). *Causality*. Cambridge university press.

Peters, R., Poulter, R., Warner, J., Beckett, N., Burch, L., and Bulpitt, C. (2008). Smoking, dementia and cognitive decline in the elderly, a systematic review. *BMC Geriatrics*, 8(1):36.

Petrylak, D. P., Tangen, C. M., Hussain, M. H., Lara Jr, P. N., Jones, J. A., Taplin, M. E., Burch, P. A., Berry, D., Moinpour, C., Kohli, M., et al. (2004). Docetaxel and estramustine compared with mitoxantrone and prednisone for advanced refractory prostate cancer. *New England Journal of Medicine*, 351(15):1513–1520.

Richardson, T. S. (2003). Markov properties for acyclic directed mixed graphs. *Scandinavian Journal of Statistics*, 30(1):145–157.

Richardson, T. S. and Robins, J. M. (2013). Single world intervention graphs (SWIGs): A unification of the counterfactual and graphical approaches to causality. *Center for the Statistics and the Social Sciences, University of Washington Series. Working Paper*, 128(30):2013.

Robins, J. M. (1986). A new approach to causal inference in mortality studies with a sustained exposure period – application to control of the healthy worker survivor effect. *Mathematical Modelling*, 7(9):1393–1512.

Robins, J. M. and Greenland, S. (1992). Identifiability and exchangeability for direct and indirect effects. *Epidemiology*, 3:143–155.

Robins, J. M. and Richardson, T. S. (2010). Alternative graphical causal models and the identification of

direct effects. *Causality and psychopathology: Finding the determinants of disorders and their cures*, pages 103–158.

Rosenbaum, P. R. (1984). The consequences of adjustment for a concomitant variable that has been affected by the treatment. *Journal of the Royal Statistical Society, Series A: General*, 147:656–666.

Rosenbaum, P. R. and Rubin, D. B. (1983). Assessing sensitivity to an unobserved binary covariate in an observational study with binary outcome. *Journal of the Royal Statistical Society, Series B: Methodological*, 45:212–218.

Roy, J., Hogan, J. W., and Marcus, B. H. (2008). Principal stratification with predictors of compliance for randomized trials with 2 active treatments. *Biostatistics*, 9(2):277–289.

Rubin, D. B. (1980). Comment. *Journal of the American Statistical Association*, 75(371):591–593.

Rubin, D. B. (1998). More powerful randomization-based p-values in double-blind trials with non-compliance. *Statistics in medicine*, 17(3):371–385.

Rubin, D. B. (2000). Comment on "causal inference without counterfactuals". *Journal of the American Statistical Association*, 95(450):435–438.

Rubin, D. B. (2006). Causal inference through potential outcomes and principal stratification: Application to studies with "censoring" due to death. *Statistical Science*, 21(3):299–309.

Satcher, D., Thompson, T., and Koplan, J. P. (2002). Women and smoking: a report of the surgeon general. *Nicotine & Tobacco Research*, 4(1):7–20.

Shepherd, B. E., Gilbert, P. B., Jemiai, Y., and Rotnitzky, A. (2006). Sensitivity analyses comparing outcomes only existing in a subset selected post-randomization, conditional on covariates, with application to HIV vaccine trials. *Biometrics*, 62(2):332–342.

Tchetgen Tchetgen, E. J. (2014). Identification and estimation of survivor average causal effects. *Statistics in Medicine*, 33(21):3601–3628.

Vansteelandt, S., Goetghebeur, E., Kenward, M. G., and Molenberghs, G. (2006). Ignorance and uncertainty regions as inferential tools in a sensitivity analysis. *Statistica Sinica*, 16(3):953–979.

Whittington, J. E. and Huppert, F. A. (1997). Smoking and cognitive decline. *Human Psychopharmacology: Clinical and Experimental*, 12(5):467–480.

Zhang, J. L. and Rubin, D. B. (2003). Estimation of causal effects via principal stratification when some outcomes are truncated by "death". *Journal of Educational and Behavioral Statistics*, 28(4):353–368.